

# Breakout 2A

## I/O Middleware, Data Management, and ties to High-Level I/O software

Session Coordinators: Lee Ward, Rob Ross

Session Scribes: Buff Miner

Session Presenter:

Session Writeup:

# Current High Level Topics

- *Reorganizing malformed accesses from applications and high-level libraries*
- *Optimizing access to underlying file systems and storage devices*
- *I/O interfaces matching HEC programming models*
- *Interfaces and models for use of intelligence by higher layers of I/O stack*
- *Alternate implementations of traditional file system services for scalability*
  - *E.g. locking, caching*

# Focus Areas

- *Small, unbalanced, unaligned accesses*
- *Collaborative caching*
- *Mapping variable sized data onto existing file systems/storage*
- *Better integration and coordination both up and down the software stack*
- *Interfaces and models for use of intelligence by higher layers of I/O stack*
- *Research to support mixed transactional/bulk I/O through the stack*
- *QoS support; common-currency, translation*
- *Robustness, improved response to errors*
- *Exploitation of RDMA and one sided operations*
- *Data format transparency*
- *Tracing*
- *Benchmarking*
- *Performance visualization*
- *Simulation*
- *Checkpointing support and optimization (e.g. deltas)*
- *Extensible metadata support*

# Consensus

- *QoS support; common-currency, translation (35/12) (long)*
- *Extensible metadata support (29/9) (long)*
- *Research to support mixed transactional/bulk I/O through the stack (27/10) (short)*
- *Robustness, improved response to errors (33/11) (continuous)*
  - *Academia doesn't see this as important*

# Contentious

- *Government emphasized*
  - Collaborative caching (23/16) (short/long)
  - Small, unbalanced, unaligned accesses (19/11) (short)
  - Data format transparency (9/9) (long)
- *Industry emphasized*
  - Interfaces and models for use of intelligence by higher layers of I/O stack (27/7) (short)
  - Checkpointing support and optimization (e.g. deltas) (18/0) (long)
- *Academia emphasized*
  - Better integration and coordination both up and down the software stack (22/3) (continuous)
  - Mapping variable sized data onto existing file systems/storage (16/3) (short)
- *Measurement and understanding (short/long)*
  - Tracing (16/0), Benchmarking (19/6), Simulation (6/1) = (41/7)
  - Matches apparent lack of interest from government in Breakout 1A
    - *Government seems only interested in benchmarking*